arXiv:1610.07004v2 [cs.CY] 26 Oct 2016

Inferring Population Preferences via Mixtures of Spatial Voting Models

Alison Nahm*, Alex Pentland†, and Peter Krafft†

*Harvard University, †Massachusetts Institute of Technology
anahm@post.harvard.edu,pentland@mit.edu,pkrafft@mit.edu

Abstract. Understanding political phenomena requires measuring the political preferences of society. We introduce a model based on mixtures of spatial voting models that infers the underlying distribution of political preferences of voters with only voting records of the population and political positions of candidates in an election. Beyond offering a cost-effective alternative to surveys, this method projects the political preferences of voters and candidates into a shared latent preference space. This projection allows us to directly compare the preferences of the two groups, which is desirable for political science but difficult with traditional survey methods. After validating the aggregated-level inferences of this model against results of related work and on simple prediction tasks, we apply the model to better understand the phenomenon of political polarization in the Texas, New York, and Ohio electorates. Taken at face value, inferences drawn from our model indicate that the electorates in these states may be less bimodal than the distribution of candidates, but that the electorates are comparatively more extreme in their variance. We conclude with a discussion of limitations of our method and potential future directions for research.

Keywords: probabilistic generative models, political polarization, demographic inference, ideal point models, computational social science

1 Introduction

Within a representative democracy, understanding the extent to which elected officials represent their constituencies is critical to evaluating the efficiency of the political system. Here we focus on the political system in the United States, where surveys typically evaluate the preferences of the electorate. Despite their widespread use, surveys can be costly, time consuming to execute, and often lack broad geographical coverage. Fortunately, there is an alternative source of readily available data about the preferences of the electorate—votes cast in elections.

The key challenge of using voting data to infer localized distributions of political preferences is the coarseness of the data. Consider inferring the distribution of political preferences of voters from votes cast in a two-candidate election. Since there are only two data points from the election (and a constraint that the two points sum to the total voting population size), inferring the distribution of preferences from these vote shares appears underdetermined.

In this paper, we introduce a model-based machine learning method to measure the political preferences of voters at a fine level of geographical granularity. To solve the underdetermination problem, we introduce a Bayesian mixture model that pools data from similar election outcomes of different geographical voting units (precincts). The method connects the distribution of preferences within each precinct to voting outcomes using a spatial voting model—a standard rational voting model from the political science literature [11,12]. Our model utilizes vote share data and a preprocessed form of campaign finance data to infer distributions of political preferences with potentially better coverage and lower cost than traditional survey methods [1,13,14,19,21].

An additional benefit of our method over surveys is that the inferred political preferences of voters are represented on the same scale as those of candidates. This is important for social science applications involving the comparison of politicians and the electorate [5]. To demonstrate the potential utility of our method and of related future work in this area, we apply our technique to understand the extent of political polarization in the Texas, New York, and Ohio electorates in comparison to that of the political candidates. While it is well-known that elected officials are highly polarized in their political positions, the political science community has not reached consensus as to whether the preferences of voters mirror this elite polarization or are comparatively moderate [1,13]. Using congressional election data for the states from the 2006, 2008, and 2010 election cycles, we find varying answers to the question depending on the polarization metric we use. We find that the distribution of the political preferences of voters is likely more extreme than that of candidates in terms of variance, while less extreme than that of candidates in terms of bimodality.

In the remainder of this paper, we begin with a discussion of related works. We then provide an overview of our novel probabilistic generative model of voting behavior. We validate this model with comparisons to results of related work and with simple prediction tasks. We then apply our model to better understand political polarization in Texas, New York, and Ohio. Finally, we conclude with a discussion of the limitations of this method and suggestions for future work.

1.1 Related Work

There has been recent work in quantitative political science that is closely related to our work. For instance, researchers recently developed a technique for estimating the preferences of the electorate and elected officials from Twitter data using a probabilistic generative network model related to the spatial voting model we use [5]. Some political scientists have used ideal point models, which are closely related to spatial voting models, to infer distributions of voter preferences from fine-grained voter data [22,16]. Unlike our work, these previous works using voting data relied on individual-level voting data, which is difficult to obtain. Other political scientists have developed meta-analysis-like methods for aggregating survey results to improve accuracy and representativeness [27], but this work still suffers from the limitation of low coverage of survey data due

to collection difficulties. Thus, the methods can only consider a coarser level of geographical granularity.

Within the computer science field, our work falls closest to a growing line of research dedicated to developing novel machine learning models for computational social science. Machine learning researchers in this area have not yet addressed the exact problem we study in our work, to the best of our knowledge. However, they have been interested in similar problems and related classes of models (e.g. [17,15,18,2]). More tangentially, a large body of work in computer science has been dedicated to drawing inferences from public observational data. Some researchers have suggested using social media data to better understand public opinion [26], while others have developed models based on inconsistent user behavior to infer their implicit preferences [10].

2 Model

We first discuss mathematical theories of voting behavior that inform our novel model. Then, we describe our model for inferring political preferences of voters.

2.1 Spatial Voting Models

Our model generalizes the widely used "spatial" or "Downsian" voting model, which is a standard model in political science of rational voting and turnout behavior in two-candidate majority vote elections [11,12]. The spatial voting model defines each voter and each political candidate as points in a one-dimensional policy space. The model defines the utility to a voter of a specific candidate winning as the Euclidean distance between their two points. Assuming the election involves exactly two candidates, the spatial voting model predicts that voters will select for their votes the candidates closest to them according to Euclidean distance in the single-dimensional policy space.

2.2 Mixtures of Spatial Voting Models

Our statistical model consists of a generative process for the vote shares of candidates in an election. Each precinct i with N_i total voters is associated with an election of exactly two candidates, c_{0i} and c_{1i}. In line with the spatial voting model, we assume that voters $v_{ji}, j \in \{1, \ldots, N_i\}$ and candidates $c_{ki}, k \in \{0, 1\}$, have positions in the same one-dimensional latent space, which we consider their *political preferences*. We assume the candidate positions are known to all election participants. Like other spatial voting models, we assume that each voter j of precinct i votes for the closest candidate in the one-dimensional latent policy space. In other words, voter j in precinct i votes for candidate 0 in precinct i if $|v_{ji} - c_{0i}| < |v_{ji} - c_{1i}|$.

We assume that each precinct is associated with a particular distribution over political preferences, which determines the preferences of the voters in that precinct. However, it is problematic to assume that these distributions are all

distinct from and independent of each other. In this case we are limited to using only one data point per precinct to infer a distinct distribution. To solve this issue we use a mixture model to pool data across precincts. We assume that certain subsets, or *clusters*, of precincts share the same distribution of preferences. These assignment of precincts to clusters is determined dynamically during inference according to similarity in observed voting patterns. This modeling assumption seems reasonable given that it is likely neighboring precincts will have similar distributions of preferences.

The expected proportion of precincts that will be assigned to each of K clusters is given a Dirichlet prior, $\boldsymbol{\theta} \sim Dirichlet(\mathbf{1})$ and $|\boldsymbol{\theta}| = K$. The assignment of each precinct i to a particular cluster is then drawn as in a standard mixture model, $x_i \sim \boldsymbol{\theta}$. The position of each voter j in each precinct i is drawn according to a component distribution associated with the cluster assignment of that precinct. The parameters of these component distributions are given weakly informative priors. The number of votes n_{ki} received by candidate k in precinct i are then given deterministically by the spatial voting model specified above. Mathematically, $n_{0i} = \sum_{j=1}^{N_i} \mathbb{1}(|v_{ji} - c_{0i}| < |v_{ji} - c_{1i}|)$, where $\mathbb{1}$ is an indicator function, and $n_{1i} = N_i - n_{0i}$.

We treat candidate positions as fixed and given since we have data on these values, but we treat voter positions, precinct assignments, and cluster distribution parameters as unknown. After marginalizing over voter positions, then conditioning on direct estimates of candidate positions and on observed vote shares per candidate, we can use Bayesian inference to arrive at likely values for the remaining unknown parameters, thus estimating the distributions of voter preferences within each precinct.

3 Data

For our empirical analysis, we use three main sources of data.

Precinct-Level Voting Results. Precincts are the finest granularity of publicly accessible aggregated vote shares. We examine congressional elections of Texas, New York, and Ohio [3]. In these cycles, Republicans won 65% of the Texas elections, and Democrats won 80% of the New York elections. We also analyze Ohio to test the ability of our method to generalize to more extreme voter distributions, as Ohio is commonly labeled by political scientists as a "swing state". We consider the election cycles 2006, 2008, and 2010 because they all depend on the same district geographic boundaries set by the 2000 U.S. Census. We omit the 2002 and 2004 election cycles to focus on recent elections. Future work could analyze longer periods.

Candidate CFscores. We incorporate quantitative estimates of the political preferences of candidates called campaign-finance scores (CFscores) [7]. CFscores are one-dimensional quantitative estimates of the political ideology of political candidates, with lower values indicating more liberal ideologies and higher values

more conservative. CFscores are recognized as effective estimates of candidate ideology when estimates for unelected candidates are needed (in contrast to DW-NOMINATE scores which only exist for winning candidates [25]).

Geographic Precinct Boundaries. We link the above-mentioned data sets with the congressional candidates running in each precinct election. We assign any precincts to the district whose geographic center fall within the specific congressional district boundary lines using the Geospatial Data Abstraction Library (GDAL/OGR) package [23,28].

4 Inference

The goal of our inference procedure is to determine likely precinct assignments ($\boldsymbol{x}$) and likely parameters of the K cluster distributions in the model described in Sect. 2. These estimates allow us to characterize the distribution of voter preferences of each precinct.

For more efficient inference, we first integrate out the voter positions ($\boldsymbol{v}$) and the precinct assignments ($\boldsymbol{x}$). All component distributions we consider allow the voter positions to be integrated out analytically. This yields the following posterior distribution:

$$P(\boldsymbol{\eta}, \boldsymbol{\theta} \,|\, \boldsymbol{v}) \tag{1}$$

$$\propto P(\boldsymbol{\theta}) P(\boldsymbol{\eta}) \sum_{\boldsymbol{x}} P(\boldsymbol{x} \,|\, \boldsymbol{\theta}) \int_{\boldsymbol{y}} P(\boldsymbol{y} \,|\, \boldsymbol{x}, \boldsymbol{\eta}) P(\boldsymbol{v} \,|\, \boldsymbol{\eta}, \boldsymbol{y}) \tag{2}$$

$$\propto P(\boldsymbol{\eta}) \prod_{i=1}^{M} \left[\sum_{x_i=1}^{K} \theta_{x_i} (\Phi_{i,x_i})^{n_{0i}} (1 - \Phi_{i,x_i})^{n_{1i}} \right] \tag{3}$$

where $\boldsymbol{\eta}$ is the distribution parameters of each cluster, and Φ_{i,x_i} is the cumulative distribution of the component distribution of precinct i after integrating out $\boldsymbol{y}$, in other words $\Phi_{i,x_i} = P\left(y < \frac{c_{1i} - c_{0i}}{2} \,|\, \eta_{x_i}\right)$. For prior distributions over the cluster parameters, when using Normal component distributions, we use a Normal prior for $\boldsymbol{\mu}$ with a mean of 0 and a variance of 100, and for $\boldsymbol{\sigma}$ we use an Inverse Gamma distribution with scale and shape parameters both set to 1.

We then use a Metropolis-Hasting Markov Chain Monte Carlo (MCMC) algorithm to arrive at likely values for the component distribution parameters and the mixture proportion ($\boldsymbol{\theta}$)[1]. To generate the results presented in this paper, we ran four independent MCMC chains from randomly generated initialized values. We then selected the set of parameter values from all generated sets that yield the highest posterior. We infer parameter values separately for the data of each state and election cycle combination described in Sect. 3. After inferring the parameters of the mixture distribution, we infer all precinct assignments to

[1] Code and data for analyses are available at `https://github.com/anahm/inferring-population-preferences`

clusters, $\boldsymbol{x}$, by selecting the maximum a posteriori (MAP) assignment variable for each precinct.

Our model-based method provides a better approach than analyzing vote shares because we can derive an overall distribution of preferences per precinct rather than a single point. This approach also has benefits even looking at coarser geographic granularity as well. We ultimately aggregate our inferences to district or state level for validation purposes, and inferring precinct-level distributions opens the possibility for the distribution at a less-granular level to be a complicated combination of precinct-level distributions.

5 Validation

To assess validity, we compare summary statistics of the distributions we infer with corresponding values from related works. In addition, we compare the predictive performance of our model with prediction methods based on empirical data and results of related works. In this section, we present results assuming the number of clusters, K, is 4 and underlying Normal component distributions, but we reach similar results when we vary the value of K and the component distribution type, which can be seen in Sect. B. These validation methods are meant to qualitatively and quantitatively assess the "face validity" of our proposed model.

5.1 Comparison with Related Works

We compare our results with the results of two survey-based methods for estimating district-level political preferences. This comparison would ensure that our inferred distributions are qualitatively reasonable from the perspective of the prior related work. A district is a coarser granularity geographical unit than a precinct, but precinct-level surveys are rarely implemented due to high costs. To compare, we obtain a single-point estimate of each congressional district preference in Texas, New York, and Ohio from our inferred precinct-level distributions.

To compute single-point estimates of district-level voter preferences, we use our model's assumption that each precinct has the same parameters as the inferred parameters of its assigned cluster. The district-level estimates are averages of the precinct-level inferences of precincts in the same district weighted by the population of those precincts.

Comparison with Raw Survey Results. We first compare with the Cooperative Congressional Election Study (CCES) [4]. The CCES surveys over 50,000 Americans every election cycle. Many political scientists use the CCES to understand the American public opinion. Moreover, the CCES respondents report their congressional districts, which yields more fine-grained data than most other national surveys [4]. We compare our results with the survey responses to two questions. The first question asks:

> Thinking about politics these days, how would you describe your own political viewpoint? (Very Liberal, Liberal, Moderate, Conservative, Very Conservative, Not sure)

As shown in Fig. 1, we find a significant positive correlation between our district-level point estimates and the responses to this question. The correlation level of the results of our method and reported survey values is 0.3127 with a p-value less than 0.01.

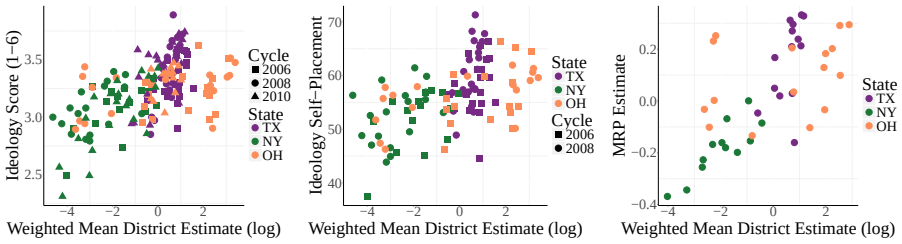

Fig. 1. (Left) Inferred district-level voter preferences compared with CCES question of self-reported ideologies on a discrete scale. (Center) Inferred preferences compared with CCES self-reported ideologies on a continuous scale from 0 to 100. (Right) District-level inferences of a decade compared with MRP estimates [27]. In all cases, the inferred district-level voter preferences are weighted mean district estimates transformed from x to $sign(x)\,log(|x|+1)$.

The second CCES question asks survey respondents to score their political ideology on a continuous scale. The question is phrased:

> One way that people talk about politics in the United States is in terms of left, right, and center, or liberal, conservative, and moderate. We would like to know how you view the parties and candidates using these terms. The scale below represents the ideological spectrum from very liberal (0) to very conservative (100). The most centrist American is exactly at the middle (50). Where would you place yourself?

This question was only used in the 2006 and 2008 CCES surveys [4]. As shown in Fig. 1, we find a significant positive correlation between our estimates and the responses to this CCES question. The correlation level of the log of the estimates of our method and reported survey values is 0.2535 with a p-value less than 0.01. Using a monotonic transformation is acceptable since the answers to survey questions and our inferred preferences are not necessarily on comparable scales.

Comparison with Aggregated Survey Results. In addition to CCES results, we validate our results against district-level ideological scores developed

by two political scientists, Chris Tausanovitch and Christopher Warshaw. They use disaggregation and multilevel regression with post-stratification (MRP) on survey data from 2000 to 2010 to estimate mean policy preferences of congressional districts [27]. Their work is one of the recent related works understanding preferences, and they analyze election cycles in a similar time frame to this paper [27]. Further, these ideological scores might be more representative of the U.S. population because the work's inference methods account for possible sampling bias.

One caveat of the work by Tausanovitch and Warshaw is that their estimates span a decade of voter behavior rather than a single election. To ensure comparison across consistent measures, we aggregate the district-level results of our model across the three election cycles into one district-level estimate spanning 2006-2010. As shown in Fig. 1, we find a significant positive correlation of 0.4149 between the log of the results of our method and their results with a p-value less than 0.01. By contrast, the MRP compared to the two CCES responses has correlations of 0.6543 and 0.7189.

The correlations between our results and the results of the survey-based methods suggest that our method can infer similar qualitative distributions to those of prior works. Further, our method can not only recover district-level mean political preferences, but also examine more granularity, precinct-level preferences in a shorter time period.

5.2 Predictive Power

As a second method of validation, we analyze the predictive power of our model. Specifically, we derive values from a *comparison election* to predict the Democratic vote share of a separate *target election*. Here a target election is either an election occurring in a later cycle or for a different government position occurring in the same cycle.

While prediction tasks are sometimes used to argue that a model has the best predictive value compared to alternative models, that is not the goal of this section. We are aware that the predictive power could be improved by incorporating more types of data. Rather, the purpose of these prediction tasks is to demonstrate that our model achieves predictive performance comparable to reasonable alternatives. Our method has additional benefits in terms over the comparison methods, so these prediction tasks are meant to lend quantitative face validity to our model.

Methods of Prediction. We compute the expected vote share of the Democratic candidate (assigned to be candidate 0) of the target election by assuming voters follow the spatial voting model and that voter preferences are identical in the comparison and target elections. The predicted vote share of each precinct i is given by Φ_{i,x_i}, the cumulative of the inferred distribution of voter preferences at the midpoint between the CFscores of the two candidates running in the target election. We aggregate vote shares of the candidate in all precincts of the same district to facilitate comparison with less fine-grained data sources.

We compare the predictive power of values yielded by our model with three alternatives: raw vote shares of previous elections, survey data, and MRP estimates. The method using raw vote shares assumes the candidates of the same political party receive the same proportion of votes in the comparison and the target election. For instance, this naïve baseline predicts the vote share of the Democratic candidate in 2010 is equivalent to the vote share of the Democratic candidate in the previous election in 2008. We consider this the baseline prediction model.

The survey prediction method uses CCES responses to a question on political party affiliation as a proxy for votes for the Democratic candidate in the target election [4]. We approximate the percentage of Democrats as the number of reported Democrats and half the reported Independent or Other divided by the total number of responses. We assume the respondents who state Independent or Other divide equally between Democrat or Republican when faced with only those options.

We also develop a prediction method based on the MRP estimates developed by Tausanovitch and Warshaw [27]. This method, labeled in Fig. 2 as MRP Cross-Val, is a simple cross-validation leave-one-out prediction method using MRP ideological scores to predict vote share. For the Cross-Val method, we obtain predictions for each district given the remaining districts and combine error terms into one mean squared error. The previously described prediction methods predict all district preferences at once, so the methods only yield one error term, which is the mean squared error term.

Prediction Tasks. We examine two prediction tasks: next cycle and same-year. The next cycle prediction task defines target elections as congressional elections of the same district one election cycle (two years) after the comparison election cycle. In other words, we use point estimates of voter preferences from election cycle $t-1$ to predict the results of election t.

The second prediction task, the same-year prediction task, defines target elections as elections for a different government position, a Senate seat, in the same election cycle as the comparison election. Although our method infers estimates using results of the same year as the election we are trying to predict, the target and comparison elections are for unrelated positions. We assume that voters consider their votes for different ballot items as in independent elections.

Prediction Task Results. As we can see in Fig. 2, in all but one case our model tends to do as well or better than the alternatives, which further validates our model. The MRP Cross-Val method is the most competitive alternative. However, the high performance of MRP Cross-Val is likely because the cross-validation method is optimized for prediction, whereas our model and the other alternatives are not.

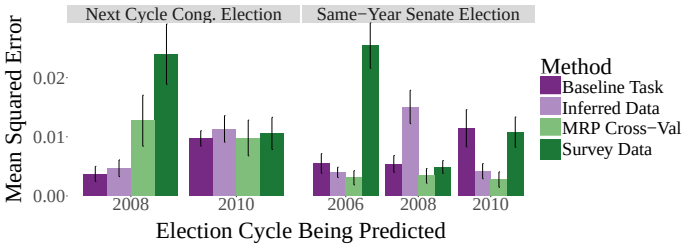

Fig. 2. Plot of the mean squared error of the actual and predicted vote share yielded by various prediction methods. Error bars are standard error of the mean.

6 Application

The results of our method distinguish themselves from the baseline methods discussed in Sect. 5 because they not only shed light on voter preferences, but also represent those preferences in the same latent policy space as known positions of political candidates. We can leverage the latter capability to assess the degree of political polarization in the electorate compared to candidates' political positions.

6.1 Background of Polarization

Previous Work. Popular media and political science communities have observed that the American political elite is becoming increasingly polarized over time, but much less work has drawn conclusions on *mass polarization* [24,29]. Some political scientists hypothesize that the distribution of the political preferences of the U.S. electorate is unimodal and moderate compared to that of the political elite [14,13]. Yet other evidence suggests increasing polarization in the American population [1]. The results of our method could add a valuable new perspective since we draw on data sources separate from the survey methods many of the opposing arguments used [1,13,14,19,21].

Among the related works that do not base their conclusions on survey data, Fiorina and Abrams find that most observations of mass polarization trends, such as differences in sociocultural attributes and world views, are not strong enough to make definitive claims about overall trends of mass polarization [13]. Their main critique of utilizing voting behavior as a proxy for political polarization is that the past actions and political preferences of candidates are not factored into the model. Our work addresses this by factoring in candidate CFscores in addition to voter behavior. Further, most related works only use one numeric metric, the difference of the means of subsets of the population, to measure mass polarization [21]. However, DiMaggio, et al. define mass polarization in terms of discrepancies between distributions, which can be described by more than a single number [9]. Some have analyzed polarization in a more holistic way, but their methods depend on survey data [19,21].

Defining Political Polarization. Mass political polarization refers to polarization in the electorate, while elite polarization refers to polarization among elected officials. One way to directly measure polarization in either case is to fit a mixture of two Normal distributions to the distribution of preferences of a population, then take the standardized difference of the component means of that mixture. For this procedure to be interpretable, we assume each mixture component has equal weight (0.5) and the same variance. We standardize the absolute difference of the two means by dividing by the inferred standard deviation of each component. This unique "difference-of-means" metric is intended as a rough proxy for the probability mass missing from the center of a distribution of preferences. Related works often use a similar metric, but they compute means by aggregating survey responses or other point estimates of subsets of the population rather than fitting a mixture model to the population distribution.

We also measure political polarization in two ways previously used in the political science literature: dispersion and bimodality [9]. Dispersion represents the extent to which more varied opinions in the population increase the difficulty for a "centrist political consensus" to exist in the population [9]. DiMaggio, et al. suggest measuring dispersion with the standard deviation of the distribution of political preferences. An increase in the standard deviation signifies that voters have more extreme political preferences and less moderate preferences in the middle of the distribution. Bimodality represents the level of separate opinions of different groups that can lead to a higher chance of social conflict. Bimodality can be measured with the kurtosis of the distribution. The formal definitions of these quantities are given in the appendix in Sect. A.

6.2 Analyzing Polarization

We apply our results to the question of trends in mass polarization for Texas, New York, and Ohio. We generate state-level distributions of voter preferences similar to the method described in 5.1. The resulting inferred distributions are shown in Fig. 3.

Mass versus Elite Polarization. Using the state-level estimates, we compute the polarization metrics as described in Sect. 6.1 for the inferred voter distributions and the corresponding candidate distributions.

First, as seen in Fig. 3, the inferred distributions of voter preferences visually appear unimodal, even though our model does not make this assumption at the state level. This is supported quantitatively in Table 1, where the electorate has relatively consistently smaller difference-of-means polarization metric than the candidates. This suggests there is no mass polarization in terms of bimodality within each state. Second, we find in Table 1 that the standard deviation of the distribution of voter preferences is generally larger than that of the distribution of candidates. This consistent difference in dispersion suggests that voters often have more extreme preferences than the candidates running for office. Third, the kurtosis of the distribution of voter preferences is generally larger than that of

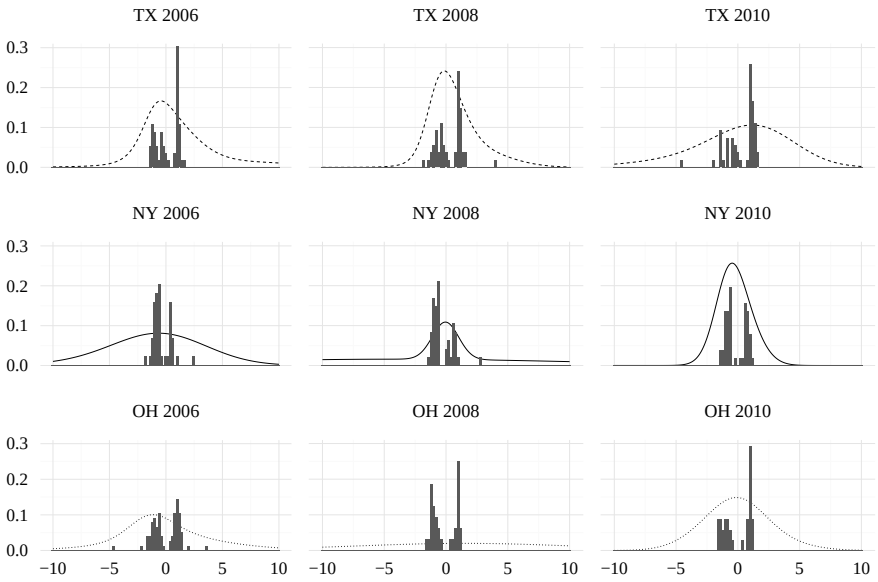

Fig. 3. Distributions of inferred voter preferences (represented by lines) and candidate preferences (histograms) based on data of the 2006, 2008 and 2010 congressional elections in Texas, New York, and Ohio.

the distribution of candidate preferences. Larger values of kurtosis indicate lower values of bimodality, so the consistent positive difference in kurtosis suggests that the voter distributions are unimodal. The results on kurtosis corroborate qualitative observation of our plots and the difference-of-means statistic.

Together these findings suggest that the distribution of political candidates may not be representative of the distribution of preferences of the electorate in the years and states we study because the extremes and the center of the voter populations are underrepresented. Further, these findings also suggest that the elites are more polarized than the electorate in terms of bimodality, but the electorate is more polarized than the elites in terms of dispersion. Inferred distributions under alternative assumed numbers of cluster components and under alternative mixture component distributions are shown in the appendix in Fig. S1 and Fig. S4. These robustness checks further support our dispersion finding but yield somewhat more mixed results on bimodality.

7 Discussion

Our methods have some limitations that could be addressed in future work. Firstly, biases inherent in our model could undermine our qualitative conclusions about polarization. Because the tails of the component distributions in our model are nearly non-identifiable, the posterior distributions in our model end

Table 1. Polarization metrics of the voters given inferred voter preferences and of the candidates based on Bonica's CFscores [8]. The "difference" of each polarization metric row is the difference between the voter metric and the candidate metric.

		2006			2008			2010		
		TX	NY	OH	TX	NY	OH	TX	NY	OH
Difference-	Voters	0.02	0.19	5.25	0.18	0.16	12.60	0.27	0.03	0.86
of-Means	Cand.	2.00	1.33	2.11	0.36	0.62	1.88	0.13	1.96	2.16
	Diff	**-1.98**	**-1.14**	**3.14**	**-0.17**	**-0.46**	**10.73**	**0.14**	**-1.94**	**-1.30**
Standard	Voters	4.36	39.54	68.20	2.14	78.46	74.23	3.88	39.30	2.75
Deviation	Cand.	0.97	0.77	1.22	1.07	0.83	0.99	1.22	0.81	1.07
	Diff	**3.38**	**38.77**	**66.97**	**1.07**	**77.64**	**73.24**	**2.66**	**38.49**	**1.68**
	Voters	3.04	20.87	2.78	1.47	19.54	0.56	0.85	16.15	0.22
Kurtosis	Cand.	-1.62	1.67	1.23	0.49	1.79	-1.77	3.24	-1.61	-1.73
	Diff	**4.66**	**19.20**	**1.55**	**0.97**	**17.75**	**2.33**	**-2.39**	**17.76**	**1.95**

up placing a large mass on preference distributions with high variances. The inference procedure could be improved to yield more reliable estimates. Similarly, our inference procedure assumes unimodal component distributions, which could bias the aggregate state level output towards unimodal distributions. As a result, we may infer high-variance unimodal precinct distributions because of the biases in the component distributions. To address these concerns, we test the ability of our procedure to recover bimodal distributions at the state level with synthetic precinct-level data. We explicitly define bimodal precinct distributions in the synthetic data, and our inference procedure was still able to correctly recover bimodal district and state-level distributions under the assumption of unimodal component distributions. This indicates that our model is able to represent and recover bimodal preference distributions even if the model is making an incorrect assumption that individual precincts are unimodal. This result on synthetic data lends some confidence that our application results on real data are not artefactual. However, our inferences on real data still have some quirks. For example, the inferred distribution of the Ohio election in 2008 is implausibly wide. In explorations of ways to address this issue, we found that adding an informative prior on the variance helped to an extent.

Another limitation of our method is that we assume the observed vote shares of elections is tied to the underlying preference distributions. This allows the model to aggregate ideological preferences of individual voters within precincts. However, some political scientists believe voters tend to vote for candidates according to party affiliation rather than ideology [6]. While this is a large limitation to our approach, future work could extend our model by estimating the extent to which people vote along party lines as opposed to according to policy preferences. As it stands, the distributions of preferences we infer can be interpreted as the distributions of *expressed* preferences of voters, if not the distribution of actual preferences.

Finally, a more fundamental limitation is that we cannot validate the shape of the inferred precinct distributions with existing survey data, since we only have observations at the district level. Future work could create a test set through an in-depth survey of the distributions of preferences of particular precincts. Validating the distributions we infer is critical to bringing our work to a level that would be useful to practitioners. At the moment we cannot tell to what extent the shape of the individual precinct distributions we infer is determined by the data versus biases from our model.

8 Conclusion

In this work, we present a model to address the problem of understanding public opinion in the context of voter preferences and political polarization. Our model builds on a long literature of spatial voting models. We use mixtures of spatial voting models to infer clusters of U.S. precincts that display similar voting patterns. Our model simultaneously infers preference distributions associated with those clusters utilizing data not only about voters, but also of candidates. These features allow us to analyze the distributions of voter preferences on their own and relative to distributions of candidate preferences. We infer voter preferences given precinct-level election results of three election cycles and three different states. We validate our inferences to the extent that we can using existing data. We validate against alternative measures of public opinion, as well as by comparing the predictive power of our model and alternative methods.

One extension of this work could adapt the model to account for elections with an uncontested candidate using similar ideas tried by Levendusky, et al. [20]. Our model could also be updated to include an offset accounting for the number of political parties in the election, which can address the bias of results based on a two-party system. Another direction is to explore other applications of our inference methods in the field of political science. For example, inferred precinct-level preference distributions could predict the effects of congressional redistricting, the process of assigning geographic boundary lines to congressional districts. Precincts are the building blocks for districts, so we could use precinct preference estimates of our model to predict the effects of redistricting proposals on the makeup of Congress.

Variations of our model could be applied more broadly beyond the scope of political science to understand distributions of preferences on other topics. For example, surveys are also used to understand consumer preferences on certain consumer products. A variation of our model could avoid the need for surveys or supplement surveys in this and other areas.

9 Acknowledgements

Special thanks to David Lazer for bringing our attention to Adam Bonica's work, to David Parkes for suggesting a reformulation of our model that enabled integrating out voter positions, and to Matt Blackwell for encouraging us to think

more about validity and identifiability. This work was supported in part by the NSF GRFP under grant #1122374. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect those of the sponsors.

References

1. Abramowitz, A.I., Saunders, K.L.: Is polarization a myth? The Journal of Politics 70(2), 542–555 (2008)
2. Airoldi, E.M., Blei, D.M., Fienberg, S.E., Xing, E.P.: Mixed membership stochastic blockmodels. In: Advances in Neural Information Processing Systems. pp. 33–40 (2009)
3. Ansolabehere, S., Palmer, M., Lee, A.: Precinct-level election data. http://hdl.handle.net/1902.1/21919 (2014)
4. Ansolabehere, S., Pettigrew, S.: Cumulative CCES Common Content (2006-2012). http://dx.doi.org/10.7910/DVN/26451 (2014)
5. Barberá, P.: Birds of the same feather tweet together: Bayesian ideal point estimation using twitter data. Political Analysis 23(1), 76–91 (2015)
6. Bartels, L.M.: Beyond the running tally: Partisan bias in political perceptions. Political Behavior 24(2), 117–150 (2002)
7. Bonica, A.: Database on ideology, money in politics, and elections: Public version 1.0. http://data.stanford.edu/dime (2013)
8. Bonica, A.: Mapping the ideological marketplace. American Journal of Political Science 58(2), 367–386 (2014)
9. DiMaggio, P., Evans, J., Bryson, B.: Have American's social attitudes become more polarized? American Journal of Sociology pp. 690–755 (1996)
10. Ding, W., Ishwar, P., Saligrama, V.: Learning mixed membership mallows models from pairwise comparisons. arXiv:1504.00757 (2015)
11. Downs, A.: An economic theory of political action in a democracy. The Journal of Political Economy pp. 135–150 (1957)
12. Enelow, J.M., Hinich, M.J.: The Spatial Theory of Voting: An Introduction. Cambridge Univ Press (1984)
13. Fiorina, M.P., Abrams, S.J.: Political polarization in the American public. Annual Review of Political Science 11, 563–588 (2008)
14. Fiorina, M.P., Abrams, S.J., Pope, J.C.: Polarization in the American public: Misconceptions and misreadings. The Journal of Politics 70(2), 556–560 (2008)
15. Flaxman, S.R., Wang, Y.X., Smola, A.J.: Who supported Obama in 2012?: Ecological inference through distribution regression. In: Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. pp. 289–298. ACM (2015)
16. Gerber, E.R., Lewis, J.B.: Beyond the median: Voter preferences, district heterogeneity, and political representation. Journal of Political Economy 112(6), 1364–1383 (2004)
17. Gerrish, S., Blei, D.M.: Predicting legislative roll calls from text. In: Proceedings of the 28th International Conference on Machine Learning. pp. 489–496 (2011)
18. Krafft, P., Moore, J., Desmarais, B., Wallach, H.M.: Topic-partitioned multinetwork embeddings. In: Advances in Neural Information Processing Systems. pp. 2807–2815 (2012)

19. Lee, J.M.: Assessing mass opinion polarization in the U.S. using relative distribution method. Social Indicators Research pp. 1–28 (2014)
20. Levendusky, M.S., Pope, J.C., Jackman, S.D.: Measuring district-level partisanship with implications for the analysis of U.S. elections. The Journal of Politics 70(3), 736–753 (2008)
21. Levendusky, M.S., Pope, J.C.: Red states vs. blue states going beyond the mean. Public Opinion Quarterly 75(2), 227–248 (2011)
22. Lewis, J.B.: Estimating voter preference distributions from individual-level voting data. Political Analysis 9(3), 275–297 (2001)
23. Lewis, J.B., DeVine, B., Pitcher, L., Martis, K.C.: Digital boundary definitions of United States congressional districts, 1789-2012. http://cdmaps.polisci.ucla.edu (2013)
24. McCarty, N., Poole, K.T., Rosenthal, H.: Polarized America: The Dance of Ideology and Unequal Riches, vol. 5. MIT Press (2006)
25. Poole, K.T., Rosenthal, H.: A spatial model for legislative roll call analysis. American Journal of Political Science pp. 357–384 (1985)
26. Ruths, D., Pfeffer, J.: Social media for large studies of behavior. Science 346(6213), 1063–1064 (2014)
27. Tausanovitch, C., Warshaw, C.: Measuring constituent policy preferences in congress, state legislatures, and cities. The Journal of Politics 75(2), 330–342 (2013)
28. United States Census Bureau: Tigerweb state-based data files: Voting districts - Census 2010. http://tigerweb.geo.census.gov (2010)
29. Zhang, Y., Friend, A., Traud, A.L., Porter, M.A., Fowler, J.H., Mucha, P.J.: Community structure in congressional cosponsorship networks. Physica A: Statistical Mechanics and its Applications 387(7), 1705–1712 (2008)

A Mathematical Definitions of DiMaggio's Polarization Metrics

Given the estimates of our model, we use the following analytical form of the standard deviation of a mixture model to measure political polarization in terms of dispersion:

$$M_\sigma = \sqrt{\left(\sum_{i=1}^{K} \frac{n_i}{\sum_{j=1}^{K} n_j}({\mu_i}^2 + {\sigma_i}^2)\right) - M_\mu} \tag{4}$$

where n_i is the total number of voters assigned to component i and M_μ is the weighted mean of the mixture distribution of voter preferences.

To measure political polarization in terms of bimodality, we use kurtosis. Kurtosis is the fourth central moment of the mixture distribution divided by the square of the variance of the mixture distribution. We use the following analytical form:

$$M_k = \frac{E[(X - M_\mu)]^4}{M_\sigma^4} \tag{5}$$

where X is a random variable drawn from the mixture distribution and hence the numerator is the fourth central moment of the mixture distribution. The

analytical form to compute the z-th central moment of the mixture distribution is below.

$$E[(X - M_\mu)^z] = \sum_{i=1}^{K} \sum_{j=1}^{z} \binom{z}{j} (\mu_i - M_\mu)^{z-j} w_i \, E[(Y_i - \mu_i)^z] \tag{6}$$

where Y_i is a random variable drawn from component i of the mixture distribution, w_i is the weight of each component, and $E[(Y_i - \mu_i)^z]$ is the z-th central moment of the ith component distribution. In our analysis, we weight each component in the mixture distribution by the proportion of the population assigned that component.

B Additional Results

In Sect. 5, we presented the results of our method assuming the underlying component distribution is Normal and the number of clusters (K) is 4. This section tests the robustness of these assumptions and presents our results when varying the underlying component distribution and the number of clusters.

B.1 Varying the Underlying Component Distribution

We test the inference procedure of our model not only assuming Normal component distributions, but also Uniform and Laplace component distributions. When assuming the distributions of voters follow a Laplace distribution, we use the same Normal prior defined for the mean of the Normal component for the location parameter and the same Inverse Gamma prior defined for the standard deviation of the Normal component for the scale parameter. When we use Uniform component distributions, we use the same Normal prior defined for the mean of the Normal component for both the minimum and the distance between the minimum and maximum parameters. The priors defined for the Normal component parameters can be found in Sect. 4. For each alternative underlying component distribution, the inferred distributions can be seen in Fig. S1, derived polarization metrics can be seen in Table S1, and prediction comparisons can be seen in Fig. S3.

Figure S2 visualizes the comparisons between the results derived of alternative component distributions and alternative data sources described in Sect. 5.1. We find significant positive correlations between our district-level point estimates and all of the alternative data sources. Assuming Laplace component distributions, our results have a correlation of 0.3216 with the responses selecting ideology given a discrete scale (left column in Fig. S2), 0.2514 with the responses selecting ideology along a continuous scale (middle column), and 0.5323 with the MRP estimates (right column). All of these correlations were significant with p-values less than 0.01. Assuming Uniform component distributions, our results have a correlation of 0.3652 with the responses selecting ideology given a discrete scale, 0.2404 with the responses selecting ideology along a continuous scale, and 0.6331 with the MRP estimates. Again, all of these correlations were significant with p-values less than 0.01.

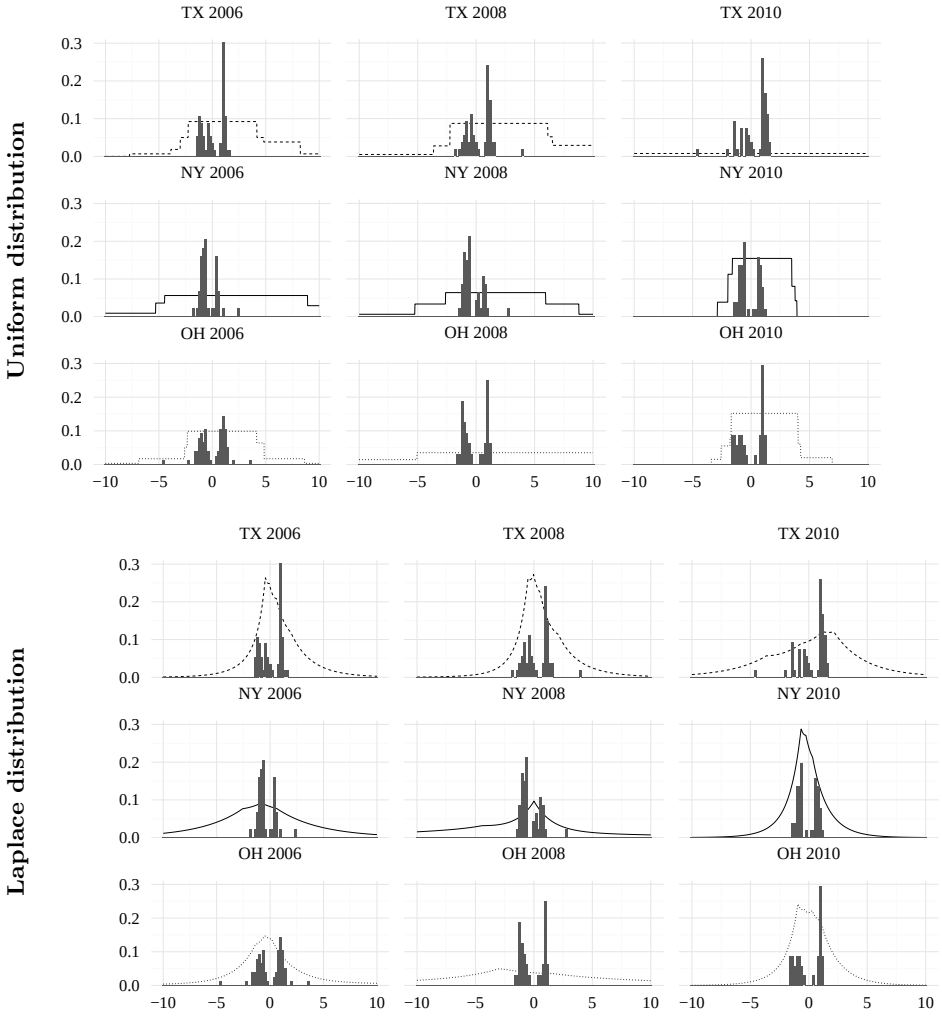

Fig. S1. Distributions of inferred voter preferences (represented by lines) under alternative assumed mixture component distributions, and candidate preferences (histograms) based on data of the 2006, 2008 and 2010 congressional elections in Texas, New York, and Ohio.

Table S1. Polarization metrics computed similarly to Tab. 1 of voters and candidates assuming different underlying component distributions.

			2006			2008			2010		
			TX	NY	OH	TX	NY	OH	TX	NY	OH
Uniform distribution	Difference-of-Means	Voters	0.02	5.73	-0.26	0.47	6.69	3.27	0.26	3.01	0.02
		Cand.	1.99	0.62	1.08	1.40	1.52	1.85	0.60	1.65	2.09
		Diff	**-1.97**	**5.11**	**-0.25**	**-0.93**	**5.17**	**1.42**	**-0.34**	**1.37**	**-2.06**
	Standard Deviation	Voters	2.11	27.51	24.55	1.94	86.17	27.71	3.42	21.46	5.12
		Cand.	0.97	0.77	1.22	1.07	0.83	0.99	1.22	0.81	1.07
		Diff	**1.13**	**26.73**	**23.33**	**0.87**	**85.34**	**26.72**	**2.20**	**20.65**	**4.05**
	Kurtosis	Voters	-0.77	42.17	-2.96	0.72	42.50	-2.99	0.38	39.39	-3.00
		Cand.	-1.62	1.67	1.23	0.49	1.79	-1.77	3.24	-1.61	-1.73
		Diff	**0.85**	**40.50**	**-4.19**	**0.23**	**40.71**	**-1.23**	**-2.86**	**41.00**	**-1.27**
Laplace distribution	Difference-of-Means	Voters	1.76	5.13	6.01	1.47	3.21	4.07	5.60	4.17	0.56
		Cand.	1.80	0.41	1.84	1.73	1.15	2.00	1.79	1.64	2.06
		Diff	**-0.04**	**4.72**	**4.17**	**-0.26**	**2.07**	**2.07**	**3.81**	**2.54**	**-1.50**
	Standard Deviation	Voters	19.25	61.64	14.51	10.64	86.42	26.79	193.24	53.56	1.67
		Cand.	0.97	0.77	1.22	1.07	0.83	0.99	1.22	0.81	1.07
		Diff	**18.28**	**60.87**	**13.29**	**9.57**	**85.60**	**25.80**	**192.03**	**52.74**	**0.60**
	Kurtosis	Voters	-2.99	1.78	1.84	-2.48	1.49	2.34	-2.90	1.47	-1.42
		Cand.	-1.62	1.67	1.23	0.49	1.79	-1.77	3.24	-1.61	-1.73
		Diff	**-1.37**	**0.11**	**0.61**	**-2.97**	**-0.30**	**4.10**	**-6.14**	**3.08**	**0.30**

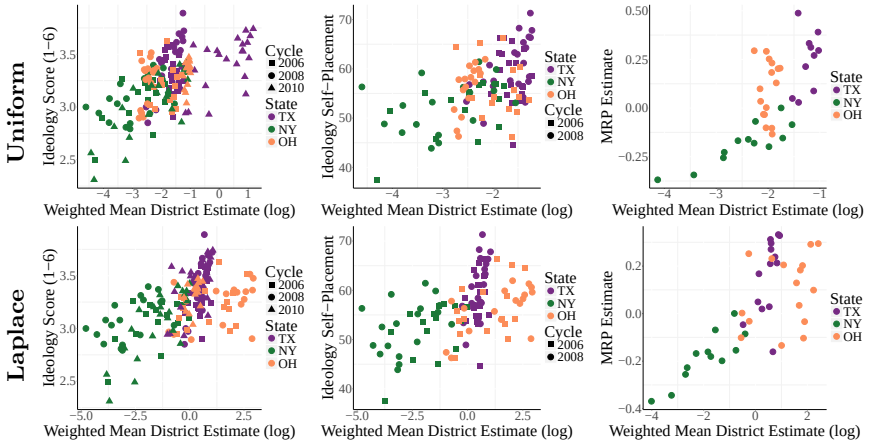

Fig. S2. In all of these cases, the inferred district-level voter preferences are based on the model varying its assumption of the underlying precinct distribution and are weighted mean district estimates transformed from x to $sign(x)\,log(|x|+1)$. (Left) Inferences compared with CCES question of self-reported ideologies on a discrete scale. (Center) Compared with the CCES self-reported ideologies on a continuous scale from 0 to 100. (Right) District-level inferences of a full decade compared with MRP estimates [27].

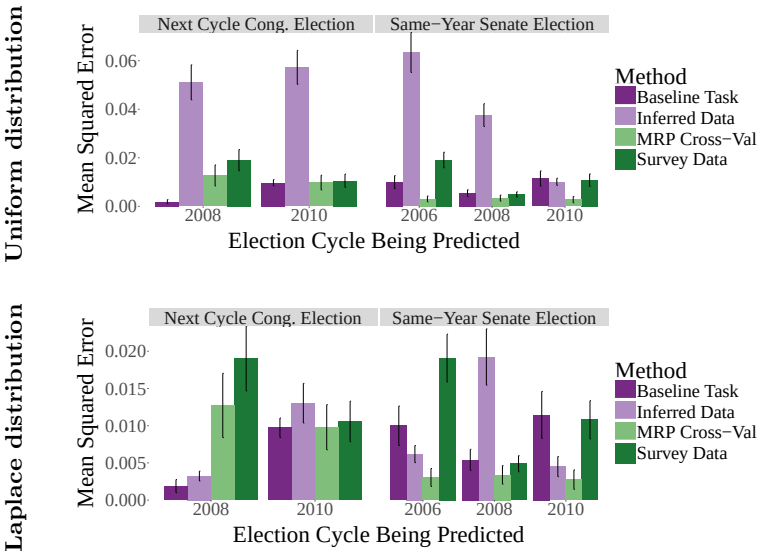

Fig. S3. Mean squared error of the actual and predicted vote share yielded by various prediction methods. Inferred data prediction method is based on our model assuming alternative underlying component distributions.

B.2 Varying the Number of Clusters

We also varied the number of clusters (K) used in the model. The main results section in the paper presented results assuming $K = 4$, but below we include the inferred distributions in Fig. S4, derived polarization metrics in Table S2, and prediction comparisons in Fig. S6 for $K = 2$ and $K = 8$, assuming Normal underlying precinct distributions. Due to time constraints, we were only able to generate these results given the Texas and New York congressional elections.

Figure S5 visualizes the comparisons between the results derived of alternative component distributions and alternative data sources described in Sec. 5.1. We find significant positive correlations between our district-level point estimates and all of the alternative data sources. When our model assumes 2 clusters rather than 4 clusters, the results of our model have a correlation of 0.2845 with the responses selecting ideology given a discrete scale (left column in Fig. S2), 0.2646 with the responses selecting ideology along a continuous scale (middle column), and 0.7160 with the MRP estimates (right column). All of these correlations were significant with p-values less than 0.01. When our model assumes 8 clusters, the results of our model have a correlation of 0.2925 with the responses selecting ideology given a discrete scale, 0.1867 with the responses selecting ideology along a continuous scale, and 0.6861 with the MRP estimates (right column). Again, all of these correlations were significant with p-values less than 0.01.

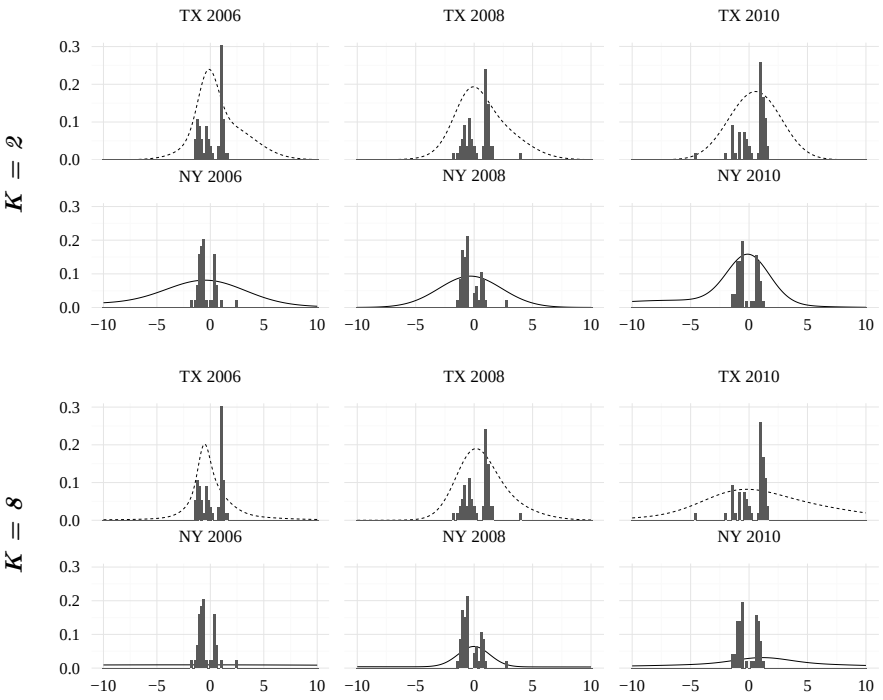

Fig. S4. Distributions of inferred voter preferences (represented by lines) under alternative assumed mixture component distributions, and candidate preferences (histograms) based on data of the 2006, 2008 and 2010 congressional elections in Texas (dotted line) and New York (solid line).

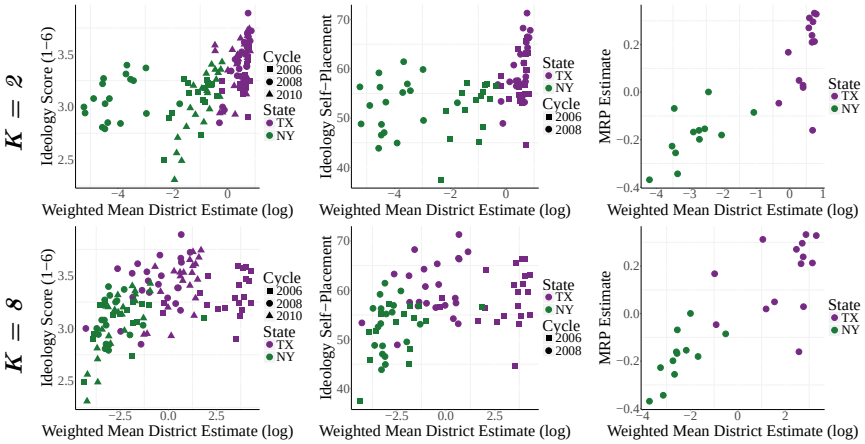

Fig. S5. In all of these cases, the inferred district-level voter preferences are based on the model varying its assumption of K and are weighted mean district estimates transformed from x to $sign(x)\,log(|x|+1)$. (Left) Inferences compared with CCES question of self-reported ideologies on a discrete scale. (Center) Compared with the CCES self-reported ideologies on a continuous scale from 0 to 100. (Right) District-level inferences of a full decade compared with MRP estimates [27].

Table S2. Polarization metrics of the voters given inferences assuming different numbers of clusters (K) and of the candidate CFscores (cand.) [8]. The "diff" of each polarization metric is the difference between the voter and candidate metric.

			2006		2008		2010	
			TX	NY	TX	NY	TX	NY
$K=2$	Difference-of-Means	Voters	0.86	2.53	0.75	2.05	1.58	1.54
		Cand.	1.87	1.58	1.83	0.37	0.41	1.433
		Diff	**-1.01**	**0.95**	**-1.08**	**1.68**	**1.17**	**0.11**
	Standard Deviation	Voters	2.40	7.75	2.36	188.98	2.15	4.56
		Cand.	0.97	0.77	1.07	0.83	1.22	0.81
		Diff	**1.42**	**6.98**	**1.28**	**188.15**	**0.94**	**3.74**
	Kurtosis	Voters	0.49	3.13	0.27	4.07	-0.15	3.24
		Cand.	-1.62	1.67	0.49	1.79	3.24	-1.61
		Diff	**2.11**	**1.46**	**-0.23**	**2.27**	**-3.39**	**4.84**
$K=8$	Difference-of-Means	Voters	9.06	9.30	3.39	6.80	1.80	1.81
		Cand.	1.74	0.54	1.82	1.55	2.17	1.42
		Diff	**7.32**	**8.76**	**1.56**	**5.25**	**-0.37**	**0.39**
	Standard Deviation	Voters	110.59	60.49	38.55	85.00	5.59	55.75
		Cand.	0.97	0.77	1.07	0.83	1.22	0.81
		Diff	**109.61**	**59.72**	**37.48**	**84.17**	**4.37**	**54.94**
	Kurtosis	Voters	5.64	2.65	48.57	4.06	-0.03	0.50
		Cand.	-1.62	1.67	0.49	1.79	3.24	-1.61
		Diff	**7.26**	**0.98**	**48.07**	**2.26**	**-3.27**	**2.11**

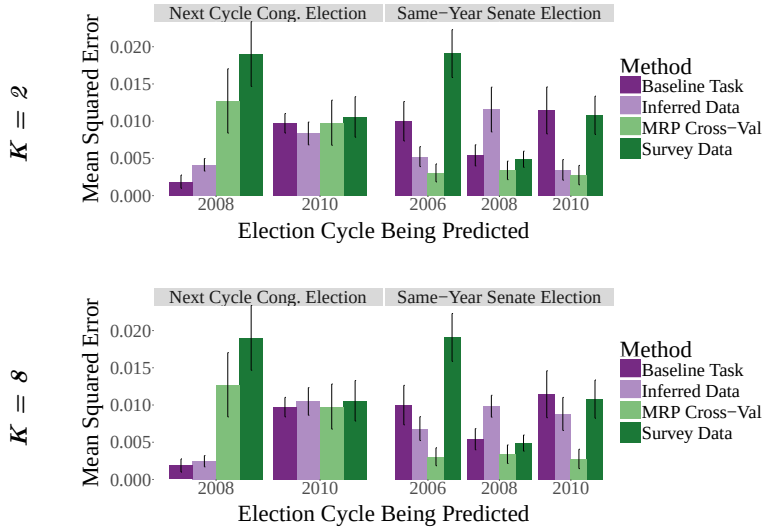

Fig. S6. Mean squared error of the actual and predicted vote share yielded by various prediction methods. Inferred data prediction method is based on our model assuming different numbers of clusters (K).